Simulating and Optimising Quantum Thermometry Using Single Photons

W.K. Tham[1,*], H. Ferretti[1], A.V. Sadashivan[1], and A.M. Steinberg[1,2]

[1]Centre for Quantum Information & Quantum Control and Institute for Optical Sciences, Department of Physics, University of Toronto, 60 St. George St, Toronto, Ontario, Canada, M5S 1A7
[2]Canadian Institute For Advanced Research, 180 Dundas St. W., Toronto, Ontario, Canada, M5G 1Z8
[*]wtham@physics.utoronto.ca

arXiv:1609.01589v1 [quant-ph] 6 Sep 2016

ABSTRACT

A classical thermometer typically works by exchanging energy with the system being measured until it comes to equilibrium, at which point the readout is related to the final energy state of the thermometer. A recent paper noted that different-temperature baths lead not only to different equilibrium states but also to different equilibration *rates*. In some cases this means that temperature discrimination is better achieved by comparing the rates than the asymptotic states – and should therefore be carried out at finite times rather than once equilibration is essentially complete. The theory work also noted that for a *quantum* thermometer, the difference between the relaxation rates for populations and coherences means that for intermediate time regimes (before full equilibration but after some characteristic time that depends on the temperatures of the baths), optimal discrimination is achieved not by probing energy only but by using quantum coherence as well. In this work, we study these effects experimentally. Implementing a recent proposal for efficiently emulating an arbitrary quantum channel, we use the quantum polarisation state of individual photons as models of "single-qubit thermometers" which evolve for a certain time in contact with a thermal bath. We investigate the optimal thermometer states for temperature discrimination, and the optimal interaction times, confirming that there is a broad regime where quantum coherence provides a significant improvement. We also discuss the more practical question of thermometers composed of a finite number of spins/qubits (greater than one), and characterize the performance of an adaptive protocol for making optimal use of all the qubits.

Introduction

A multitude of measurement and metrology tasks have been shown to benefit, sometimes dramatically, from the substitution of certain classical resources with their quantum counterpart[1,2]. Examples of advances in quantum metrology range from many-fold increase in sensitivity to phase in interferometry or polarization in polarimetry through the use of non-classical light[3–10] to highly sensitive magnetometry[11–13]. Much in keeping with the spirit of the field, a recent analysis[14] has considered that most primitive of metrological tasks - thermometry, or simply telling cold from hot - and found that a coherent measurement scheme can enhance thermometry beyond the traditional approach of allowing a measurement device to fully thermalize and equilibrate with the thermal bath being measured.

Our measurement device in this case is a qubit or a quantum system consisting of two levels typically called the ground (denoted $|0\rangle$) and excited ($|1\rangle$) states. Furthermore, since a single spin cannot provide more than one bit of information[15], we follow[14] in considering the simpler task of distinguishing between just *two* (instead of a continuum of) thermal baths at two different temperatures. And whereas a traditional thermometry approach prescribes letting the thermometer fully thermalise to the given bath by waiting for a long time, whereupon some physical quantity that bears a known correlation with the temperature is measured, we allow our qubit to interact for only a finite duration before it is subjected to some measurement. The qubit interacts with the heatbath by absorbing or emitting photons from/into it. In standard treatments of spontaneous emission[16,17] this photon exchange process is often described with parameters τ_1, the decay time for the excited state population, and τ_2, the decay time for coherences between $|0\rangle$ and $|1\rangle$. It is known that $2\tau_1 \geq \tau_2$. When there are no additional dephasing mechanisms, the equality holds so that the coherence damps away exactly half as quickly as the excited state population. In the case of thermalisation via energy exchange with a bosonic bath, as treated in[14], the collision rate grows with occupation number, and hence with temperature. τ_1 and τ_2 are therefore shorter for higher-temperature baths. This difference means that it is generally advantageous to wait for a finite interaction time (on the order of the thermalisation times) in order to better distinguish the two processes, rather than allowing the qubit to fully thermalise with the bath. In particular, for high temperatures, the asymptotic populations are very similar, while the thermalisation rates may be quite different. Since coherences decay twice as slowly as populations, the time of optimal distinguishability occurs roughly twice as late for initially coherent states (which decay principally at τ_2) as for initially incoherent states (which decay at τ_1). As a consequence, one can show that beyond a

certain critical time, an initial state with some coherence always makes for a more sensitive thermometer. Similar results are expected to hold in any case where thermalisation occurs faster with higher-temperature baths (for instance, via collisions with Maxwell-Boltzmann distributed gas molecules), but not if the thermalisation rate were fixed and temperature-independent, as in commonly used phenomenological models of thermal conductance. (For the interested reader, we give a mathematical argument in the Supplementary, as to why these observations are true). This work therefore aims to experimentally demonstrate this metrological advantage, along with an extension to the more practically relevant case where one is not restricted to the use of a single qubit.

In the Bloch sphere representation, every single qubit state corresponds uniquely to a 3-vector on or within the unit sphere. Conventionally, the excited state is represented as $+Z$ and the ground state as $-Z$. Maximally coherent states lie on the xy-plane, usually with $(|0\rangle + |1\rangle)/\sqrt{2}$ at $+X$. The thermalisation of the qubit in this picture can be thought of as a trajectory from its initial Bloch vector (usually a unit vector on the surface of the unit sphere if the initially prepared state is pure) to its final point on the z-axis. To aid the reader in visualising this, figure 1 shows trajectories of a qubit initialised in $+Z$, $-Z$, and $+X$. For the $\pm Z$ initial states, the evolution of the qubit is strictly along the z-axis, so only the z component (denoted s_z) is shown. When the qubit is initialised to $+X$, however, both the coherence (the x-component, labeled s_x) and the excited state population (the z-component) relax with time so both are shown in figure 1(c) over a range of discrete times.

The yardstick by which we will characterise the performance of our qubit thermometer is the probability with which we mis-identify our bath, either by mistaking a hot bath for a cold one or vice versa. Since a thermalised qubit is almost always in a mixed state (a statistical mixture of pure states), this error probability or p_e never vanishes. For example, a fully thermalized qubit at temperature $T = \infty$ is in a state that is an equal statistical mixture between $|0\rangle$ and $|1\rangle$ whereas at $T = 0$ it is in the ground or $|0\rangle$ state. Suppose we now identify our heatbath as the hot one if and only if a measurement on the qubit finds it in the excited state. Although we will never misidentify the $T = 0$ bath, the $T = \infty$ bath yields the excited state with 50% probability so we stand to misidentify it half the time! Assuming that a given bath is chosen from $T = \infty$ and $T = 0$ with equal likelihood, our overall error probability, p_e, is 1/4. Tasks such as the one just described are aptly called state discrimination. Conveniently, p_e in state discrimination is well-known to be related to $\vec{r}_1$ and $\vec{r}_2$, Bloch vectors corresponding to the states being discriminated, as follows[18,19]:

$$1 - p_e \leq \frac{1}{2}\left(1 + \frac{1}{2}|\vec{r}_1 - \vec{r}_2|\right) \tag{1}$$

where the norm is to be understood as the usual Euclidean/Cartesian distance between vectors $\vec{r}_1$ and $\vec{r}_2$. Returning to our example with $T_{hot} = \infty$ and $T_{cold} = 0$, the completely mixed state corresponds to $\vec{r}_1 = \langle 0,0,0\rangle$ and the ground state to $\vec{r}_2 = \langle 0,0,-1\rangle$, so $|\vec{r}_1 - \vec{r}_2| = 1$ implying once more that $p_e \geq 1/4$. It is important to stress that while a larger Euclidean distance between Bloch vectors implies that a lower error probability is achievable in principle, actually saturating the inequality to achieve the lowest possible error requires that we select the correct basis during measurement. In our example above, measuring if the qubit is in the state $(|0\rangle + |1\rangle)/\sqrt{2}$ would have yielded a "yes" answer with probability 50% for *both* $T = 0$ and $T = \infty$, giving us no information at all about the bath!

Experimental Design

Emulating thermalisation with photons

In designing a tabletop experiment to test the use of coherence and adaptivity in qubit thermometry, we have opted to optically simulate the heatbaths instead of subjecting our qubits to actual thermalization. To do so, we must first understand how thermalisation affects qubits. In the absence of extraneous damping processes (e.g. mechanical collision), a qubit that interacts with a thermal reservoir by photon exchange alone can be treated as a system that emits or absorbs a photon into/from the reservoir with some probability. Such a process is well-modeled by a generalised amplitude damping (GAD) channel, which is defined in the standard operator-sum representation as follows[19]:

$$\rho_{thermalized} = p\left(K_{11}\rho_{initial}K_{11}^\dagger + K_{12}\rho_{initial}K_{12}^\dagger\right) + (1-p)\left(K_{21}\rho_{initial}K_{21}^\dagger + K_{22}\rho_{initial}K_{22}^\dagger\right) \tag{2}$$

$$K_{11} = \begin{bmatrix} 1 & 0 \\ 0 & \sqrt{1-\gamma} \end{bmatrix} \qquad K_{21} = \begin{bmatrix} \sqrt{1-\gamma} & 0 \\ 0 & 1 \end{bmatrix}$$
$$K_{12} = \begin{bmatrix} 0 & \sqrt{\gamma} \\ 0 & 0 \end{bmatrix} \qquad K_{22} = \begin{bmatrix} 0 & 0 \\ \sqrt{\gamma} & 0 \end{bmatrix} \tag{3}$$

Equivalently we can write the above process in terms of its action on the Bloch vector:

$$\vec{v}_{thermalized} = \begin{bmatrix} \sqrt{1-\gamma} & 0 & 0 \\ 0 & \sqrt{1-\gamma} & 0 \\ 0 & 0 & (1-\gamma) \end{bmatrix} \vec{v}_{initial} + \begin{bmatrix} 0 \\ 0 \\ (2p-1)\,\gamma \end{bmatrix} \tag{4}$$

In words, the Kraus operators describe two physical sub-processes: K_{11} and K_{12} jointly describe a sub-process in which the qubit in $|0\rangle$ *absorbs* a photon from the reservoir with probability γ thereby transitioning to $|1\rangle$. K_{21} and K_{22} describe the opposite sub-process in which a photon is *emitted* into the reservoir again with probability γ. A thermalising qubit is merely one which undergoes the first sub-process (absorption) with probability p and the second (emission) with probability $1-p$. The probability p in turn is determined by the bath temperature. To see this, suppose our qubit states $|0\rangle$ and $|1\rangle$ have energies E_0 and E_1 respectively. After fully thermalising, we expect our qubit, which is now in a mixture of $p\,|1\rangle$ and $(1-p)\,|0\rangle$, to obey thermal statistics. We expect:

$$\begin{aligned} p &= \frac{\exp(-E_1/k_BT)}{\exp(-E_0/k_BT)+\exp(-E_1/k_BT)} \\ \implies \ln\frac{1-p}{p} &= \frac{E_1-E_0}{k_BT} = \frac{\hbar\omega}{k_BT} \end{aligned}$$

so $p \to 0$ as $T \to 0$ whereas $p \to 1/2$ as $T \to \infty$. It is also well-known that for a bosonic thermal reservoir, the Planck distribution implies an average occupation number, $\overline{N} = [\exp(\hbar\omega/k_BT)-1]^{-1}$. Thus, in terms of $\overline{N}$, we can write p more compactly as $p = \overline{N}/\left(1+2\overline{N}\right)$.

The damping parameter γ is related to the interaction time with the bath, t, and temperature, T, as follows: $\gamma = 1 - \exp(-t/\xi\,\tau_{sp})$, where τ_{sp} is a timescale characteristic of the coupling between qubit and bath. Here,

$$\xi = \tanh(\hbar\omega/2k_BT) = \frac{1}{1+2\overline{N}}$$

is a unitless quantity that encodes the bath temperature. Note from equation 5 that ξ, and therefore the relaxation rate of the qubit's excited state population and coherence, is temperature dependent. This, coupled with the fact that the difference of two exponential functions with different exponents is *not* monotonic, means that the Euclidean distance between resultant states for the hot vs cold baths is larger (and our thermometer more sensitive) when the interaction time t is finite (partially thermalised) as opposed to infinite (fully thermalised).

Since in our case we are merely emulating thermalisation, ω which defines the mode through which the qubit is coupled to the bath, is ill-defined. We shall therefore specify temperature in terms of $\overline{N}$ and $\hbar\omega/k_B$. Likewise, we do not have an intrinsic timescale by which to specify τ_{sp}. All times t will be specified in units of τ_{sp}. As we'll see, encoding the qubit in two orthogonal polarisations of a photon (call them $|H\rangle$ and $|V\rangle$) makes it possible to implement this compactly in a tabletop experiment.

Implementation

We use a type-II spontaneous parametric down-conversion (SPDC) setup with a $2mm$ BBO crystal pumped by a $50mW$ continuous wave $405nm$ diode, phase-matched with an opening angle of 3°, and filtered with 10nm bandpass filters. One photon of the SPDC pair is sent directly to a single-photon counting module (SPCM) to act as herald whereas the other is used as a $810nm$ light source for our experiment. Typical total coincidence rates (taking into account all losses in the experiment) are ≈ 4000 counts per second. The first stage of the experiment immediately following the SPDC source is a polariser $\rightarrow$ quarter-wave plate (QWP) $\rightarrow$ half-wave plate (HWP) sequence, which collectively comprises the "state preparation" block in figure 2.

To simulate thermalization, we make use of an optical circuit having the capacity to simulate any valid single-qubit quantum channel (precisely, completely positive trace-preserving or CPTP maps). The design of this circuit was inspired by Sanders et al's theoretical results[20], which were in turn based on mathematical work by Ruskai et al[21]. A similar optical circuit has already been shown to accurately simulate a wide variety of single-qubit quantum channels[22]. Our design, illustrated in figure 2 consists of a variable beamsplitter (VBS), followed by two subsequent interferometers (labeled channels 1 and 2 respectively in figure 2). The optical design of each channel is shown in inset (b) of figure 2. Light that is incident on each channel is split at a polarising beamsplitter (PBS) so that each of two orthogonal polarisations travels along spatially separated counter-propagating paths. Each path contains a HWP, allowing the polarisations to be rotated independently before being recombined at the PBS. It is easy to see how this intra-interferometer rotation allows us to emulate a damping channel. Suppose we leave the $|H\rangle$ polarisation within the interferometer unrotated but leave the HWP within the $|V\rangle$ path at 45° so that $|V\rangle \rightarrow |H\rangle$. Such a setting guarantees that the output of the channel is *always* $|H\rangle$ regardless of input polarisation - hence a full damping channel with

$\gamma = 1$. Less extreme settings realise the full range of damping channels. More formally, it is easy to show that each channel is well described by $\rho_{output} = K_1 \rho_{in} K_1^\dagger + K_2 \rho_{in} K_2^\dagger$ with:

$$K_1 = \begin{bmatrix} \cos 2\theta_H & 0 \\ 0 & \cos 2\theta_V \end{bmatrix} \qquad K_2 = \begin{bmatrix} 0 & \sin 2\theta_V \\ \sin 2\theta_H & 0 \end{bmatrix}$$

Here θ_H and θ_V are angles of the fast axis of the half-wave plates in channels 1 and 2, acting on the horizontal and vertical paths respectively. Setting $\theta_H = 0$ and $\sin 2\theta_V = \sqrt{\gamma}$ implements K_{11} and K_{12} in equation 3, whereas setting $\theta_V = 0$ and $\sin 2\theta_H = \sqrt{\gamma}$ implements K_{21} and K_{22}.

Although the inner workings of the VBS are not shown, it is identical in design to the channels except for the fact that there is a single HWP (instead of two independent ones) acting on both counter-propagating paths. This restricts the action of the VBS to a *fixed* unitary that can be subsequently undone for all input states with a simple HWP placed outside the VBS. Notice that if we now set the HWP in the VBS to $\cos^2 2\theta_{VBS} = p$, both channels work in tandem to fully implement the map in equation 2. Thus, we have a fully tunable means of emulating thermalisation with a polarisation qubit.

Finally, the measurement block consists of a QWP $\rightarrow$ HWP $\rightarrow$ PBS sequence. The four output states (two from each channel, one from each Kraus operator) are mixed and sent through this measurement sequence and then onto an APD and coincidence counter. In practice, losses and imperfections in optical components means that mixing the four output states *before* the measurement sequence is impractical since it precludes the possibility of compensating with post-processing. Instead, we opted to send each one to two APDs after the measurement sequence (they are further split at the PBS into $|H\rangle$ and $|V\rangle$) and then only tracing over them in post-processing. Because the number of settings for state preparation, measurement, and channel selection is potentially vast, the HWPs and QWPs are mounted on motorised rotation mounts (Thorlabs PRM1/MZ8) driven by DC servo controllers (Thorlabs KPRM1E/M) where necessary. In addition to automating time consuming parts of the experiment, they provide the added benefit of more precise angular settings ($\pm 0.2°$ as opposed to roughly $\pm 0.5°$ when done manually).

Experimental Data

To ensure that our channels are emulating the desired thermalisation process, we characterise it via full process tomography[23–26]. Throughout stretches of data-taking, we re-characterise periodically (approximately every 30 minutes, the minimum time within which visibility of interferometers are likely to have dropped appreciably) and realign optics as necessary. We proceeded to prepare states $-Z$ (or $|H\rangle$), $+Z$ (or $|V\rangle$), and $+X$ (or $|D\rangle = (|H\rangle + |V\rangle)/\sqrt{2}$). We emulate the heat baths specified in[14], with $\xi_{cold} = 1/\left(1+2\overline{N}\right) = 1/12$ and $\xi_{hot} = 1/20$. These correspond to temperatures of $\approx 6\hbar\omega/k_B$ and $10\hbar\omega/k_B$ respectively. We further set the channels to emulate interaction times ranging from $t = 0\tau_{sp}$ to $t = 0.4\tau_{sp}$, where asymptotics have yet to dominate and dynamics are non-trivial. The time steps are sampled unevenly because: a) at large t's, we were limited by the precision of our motorised rotation stage ($\sim 0.2°$) whereas b) at small t's, we restricted ourselves to $\Delta t \geq 0.2\tau_{sp}$ in order to maintain a reasonable sampling stepsize. For each channel setting, we chose to measure along a basis prescribed by static state discrimination strategies to be optimal (i.e. select a projector M s.t. equation 6 is minimized, discussed in next section). We counted photons for 10 seconds per measurement, which yielded measurement "shots" that consist of $\sim 40,000$ photons apiece. Although the thermometry scheme discussed above is intended for single qubits, we have opted to use a bright source and long count durations in order to infer $\mathrm{Tr}(\rho M)$, the probability of a successful outcome of the projector, M, in our chosen measurement basis. Since we expect the number of coincidences to be binomially distributed, this inferred detection probability tends to the true single-photon probability with diminishing uncertainty as the total photon number becomes large. Figures 3 show this inferred probability. Note that each point in the plots represents an average over many sets of data taken under identical experimental conditions (9 sets for the $+Z$ case, 10 sets for $+X$, and 4 for $-Z$, where a hardware issue forced us to discard 5 sets of data; discarded sets are reported in the Supplementary, for completeness).

In order to compare with theory, the discrimination error probability p_e can be computed from the above detection probabilities as:

$$p_e = \min\left\{0.5\mathrm{Tr}(\rho_1 M) + 0.5\left(1 - \mathrm{Tr}(\rho_2 M)\right), 0.5\left(1 - \mathrm{Tr}(\rho_1 M)\right) + 0.5\mathrm{Tr}(\rho_2 M)\right\} \tag{5}$$

where $M = |\theta\rangle\langle\theta|$, $|\theta\rangle = \cos\theta|H\rangle + \sin\theta|V\rangle$, and $\mathrm{Tr}(\rho M)$ are detection probabilities shown in figure 3. The results are shown in figure 4. Overlayed are theory curves deduced (via equation 1) from the Euclidean distances between final states of the ideal GAD. Reiterating theoretical results mentioned above, we see that after approximately two thermalisation times ($t \sim 0.1\tau_{sp}$) the coherent state $+X$ outperforms the incoherent ones $\pm Z$. Note that although $-Z$ appears to be *globally* optimal

(i.e. has a lower p_e at $t \sim 0.07\tau_{sp}$ than other states at any time), this is peculiar to our choice of temperatures and is not always the case.

While the behaviour of p_e shows good agreement with theory for $\pm Z$ input states, the $+X$ case shows discrepencies in the region $t \leq 0.2\tau_{sp}$. This can be ascribed to the fact that our interferometers have finite visibility. This has the effect of mapping some amount of coherence between $|H\rangle$ and $|V\rangle$ to an incoherent mixture and is completely analogous to extraneous dephasing processes (e.g. atomic/molecular collisions etc.) that we did not consider in our thermalisation model. Note that $\pm Z$ input states do *not* experience interference effects as they traverse the channels - these states end up traveling through the interferometer via just *one* of the two possible counter-propagating paths and have nothing to interfere with when they re-emerge at the PBS. The $\pm Z$ states are therefore not susceptible to imperfect interferometer visibility. The same cannot be said of the $+X$ state. The interferometers that comprise our channels have typical visibilities $\geq 95\%$, though due to the lack of active stabilization, realignment can become necessary from time to time.

While the notion of the Euclidean distance and the error probability are the correct figures of merit to use for state discrimination tasks, the former applies strictly to single-qubit states whereas the latter becomes increasingly difficult to compute for large numbers of qubits. An alternate measure is the distinguishability, often used as a measure of the ease with which two distributions can be distinguished. It is defined as the squared difference of the means of the two distributions, divided by their variance:

$$\text{distinguishability} = \frac{(\text{difference in mean})^2}{\text{variance}} = \frac{[E(P_{hot}) - E(P_{cold})]^2}{\max\{\text{Var}(P_{hot}), \text{Var}(P_{cold})\}}$$

where P_{hot} and P_{cold} are the binomially distributed outcomes of some projector observable given the output state from each of the heat baths (i.e. the probabilities shown in figure 3). A plot of this measure is shown in figure 5. We attribute the noise in the experimental points in this figure to the fact that relatively few sets of data (9 sets for +Z and +X inputs and 4 sets for -Z) were used to infer the variances. Although qualitatively quite similar to the p_e plot in figure 4 (i.e. the $+X$ state remains optimal after some time, while $-Z$ is optimal if time is not a constraint), the two measures disagree for example on the cutoff times at which the optimal input state changes. A clear advantage for $+X$ is seen at $t > 0.08\tau_{sp}$ (from theory curves in figure 5), whereas in the Euclidean distance or single-qubit p_e case it is seen after $t > 0.1\tau_{sp}$. This discrepancy leads one to suspect that the optimal input state for our thermometer in a multi-qubit scenario is different from the single-qubit case. In order to obtain a more rigorous measure of many-qubit distinguishability, we numerically computed the error probability p_e for 100 qubits. This is shown in figure 6 along with the fidelity, $\mathscr{F} = \text{Tr}\sqrt{\sqrt{\rho_{hot}}\rho_{cold}\sqrt{\rho_{hot}}}$ where ρ_{hot} and ρ_{cold} are final states from the hot and cold baths respectively. The quantity $(1-\mathscr{F})/2$ has been shown[18] to bound p_e from above in the limit of asymptotically large number of qubits. Again, the differences between a single qubit vs many qubits is clear - the crossover between $-Z$ and $+X$ in figure 6(b) occurs at $t = 0.0828\tau_{sp}$. In the next section, we treat the many qubit case more carefully.

Multi-qubit extension and adaptive state discrimination

A Bayesian approach

While a single-qubit thermometer is conceptually interesting, it is obviously of more practical relevance to consider a thermometer composed of many spins, but potentially a limited, fixed number. Already, the static strategy - doing the same thing on all N copies of a qubit - yields better binomial statistics the larger N is. Let's return to the example in the previous section, of distinguishing $T = 0$ from $T = \infty$. The best static strategy for N qubits in that case is to measure along some optimal axis for all qubits and concluding that the bath is at $T = 0$ if and only if *all* measurement outcomes are 0. In that case, $p_e = {p_{outcome0}}^{N+1}$.

One can often do better[27,28] by allowing for an adaptive strategy (say by changing the measurement basis for each qubit), a possibility that was investigated by Wiseman et al[29]. To facilitate further discussion, we now restate the discrimination problem in slightly more formal terms. Whereas in the preceeding discussion we assumed that the bath was equally likely to be in T_{hot} or T_{cold}, we now allow each to occur with different prior probabilities (call them π_1 and π_2 where $\pi_1 + \pi_2 = 1$). We also assume that for a given interaction time t, the output states from baths T_{hot} and T_{cold} are known to be ρ_1 and ρ_2 respectively. The problem is to find an optimal *strategy* that yields $p_{e,min} = \min_{strat}\{p_{e,strat}\}$, the lowest error over all possible strategies. A *strategy*, in turn, is a combination of a measurement (observable M) and a *threshold* ($k \in Z$) beyond which one concludes that the output state is ρ_1 (or ρ_2). In this language, the error probability for each case (assuming we employ a strategy non-adaptively, i.e. the same M for all N qubits):

$$\begin{aligned} p_{e,1} &= \pi_1 \sum_{j=0}^{k-1} \frac{N!}{j!(N-j)!} [\text{Tr}(\rho_1 M)]^j [1-\text{Tr}(\rho_1 M)]^{N-j} + \pi_2 \sum_{j=k}^{N} \frac{N!}{j!(N-j)!} [\text{Tr}(\rho_2 M)]^j [1-\text{Tr}(\rho_2 M)]^{N-j} \\ p_{e,2} &= \pi_2 \sum_{j=0}^{k-1} \frac{N!}{j!(N-j)!} [\text{Tr}(\rho_2 M)]^j [1-\text{Tr}(\rho_2 M)]^{N-j} + \pi_1 \sum_{j=k}^{N} \frac{N!}{j!(N-j)!} [\text{Tr}(\rho_1 M)]^j [1-\text{Tr}(\rho_1 M)]^{N-j} \end{aligned} \quad (6)$$

The optimal error probability is therefore computed as $p_{e,min} = \min_{k,M}\{p_{e,1}, p_{e,2}\}$. If ρ_1 and ρ_2 are two single-qubit states, the minimization over M reduces to a minimization over one real parameter (i.e. the measurement angle θ; we can assume ρ_1 and ρ_2 both lie on the real plane of the Bloch sphere - else we rotate them onto it).

In the adaptive case where θ may vary with each copy of ρ, the above no longer holds. Instead, consider the following approach: suppose on the first measurement M_1, we obtain a successful ($\checkmark$) outcome. We update our state of knowledge as follows:

$$p(\rho_1|\checkmark) = \frac{\pi_1 p(\checkmark|\rho_1)}{p(\checkmark)} = \frac{\pi_1 \mathrm{Tr}(\rho_1 M_1)}{\pi_1 \mathrm{Tr}(\rho_1 M_1) + \pi_2 \mathrm{Tr}(\rho_2 M_1)}$$

$$p(\rho_2|\checkmark) = \frac{\pi_2 p(\checkmark|\rho_2)}{p(\checkmark)} = \frac{\pi_2 \mathrm{Tr}(\rho_2 M_1)}{\pi_1 \mathrm{Tr}(\rho_1 M_1) + \pi_2 \mathrm{Tr}(\rho_2 M_1)} \tag{7}$$

We then take these as the "updated" priors: $\pi_1^{(1)} = p(\rho_1|\checkmark)$ and $\pi_2^{(1)} = p(\rho_2|\checkmark)$ where the superscript indicates that they are post-measurement-1. Had the outcome of M_1 been negative, we would simply have replaced every instance of $\mathrm{Tr}(\rho M)$ in equation 7 with $1 - \mathrm{Tr}(\rho M)$. We then return to equation 6 to minimize θ substituting $\pi_1 \to \pi_1^{(1)}$ and $\pi_2 \to \pi_2^{(1)}$. Now armed with M_2, we proceed to repeat the procedure on the next copy of ρ. After all qubits are measured, we finally choose the state associated with the larger of $\left\{\pi_1^{(n)}, \pi_2^{(n)}\right\}$ as the output of our discrimination procedure. In[29], it was shown that such a Bayesian prior update strategy is not only better than the static one but is in fact optimal if ρ_1 and ρ_2 are pure states! And although it isn't the optimal strategy for an arbitrary mixed state, is nevertheless performs better than the static approach.

Testing Adaptivity

To test the adaptive approach, we set the channel to identity (it simply preserves all input states; in practice this meant setting $p = 0$ in the VBS and $\theta_H = \theta_V = 0°$ for the channels). The states to discriminate were $|H\rangle$ and $|D\rangle$. These states were chosen (instead of outputs of the thermalisation channel described in preceeding sections) because the identity channel is much easier to control, allowing us to avoid potential confounding factors when demonstrating the benefits of adaptivity. The identity channel requires only one of channels 1 and 2 to be active and so does *not* require LCWP calibration or careful alignment of the VBS, among other things.

As before, we began by validating the channel via process tomography and then proceeded to measure along bases specified by equations 6 and 7 above. In an actual adaptive scenario, one ideally updates one's measurement setting as each photon is detected. However, our motorised waveplate mounts are relatively slow making it difficult for us to adapt our measurement setting conditioned on individual detection events. Instead, we have chosen to measure along *all* bases that are prescribed by our Bayesian update strategy given *all* possible detection outcomes. This allowed us to map out the full tree of possible outcomes along with the probability of occurrence for each node. Such a map of outcome probabilities allows us to confirm that the Bayesian update rule is valid, even if we can't directly emulate adaptivity. Figure 7 shows a plot of how the error probability, p_e, scales with number of qubits in the various scenarios.

Figure 7 shows the resulting error probabilities, p_e, for various strategies. Theoretical predictions and experimentally derived values for our strategy, detailed above, are plotted in blue and red, and are labelled "Adaptive". For comparison, theory predictions for two static (non-adaptive) cases are shown. The "1-qubit optimum" naively uses a measurement basis that is optimal for just 1 qubit, and repeats it as necessary. On the other hand, in the "Global Optimum" approach, one is assumed to have been told the total number of qubits available, and a *static* measurement angle that is optimal for the given number of qubits is computed and used.

Evidently, the *adaptive* multi-qubit scenario offers the benefit of lower error probabilities. The absolute reduction is particularly pronounced for the first few additional qubits, when p_e is still relatively large. The deviation between our data and theoretical prediction for p_e becomes pronounced as the number of qubits N becomes large. The probability tree that we must reconstruct grows quickly with N (generally, 2^N branches) and so does the precision with which we must set our measurement basis in order to maintain an advantage over the non-adaptive approach. In our case, we are limited by our motorised waveplate mounts to an angular precision no better than $\pm 0.2°$.

Also of note, our adaptive strategy makes no assuption about the total number of qubits and continues to work in a "rolling" fashion even if, midway through the scheme, we were suddenly told that more qubits have suddenly become available. This is in stark contrast to the "global optimum" static strategy, where a favorable scaling is only possible given full knowledge of just how many qubits there are.

Conclusions

In summary, we have simulated the thermal equilibration of a spin by using a construction of a universal emulator for quantum channels. This has allowed us to confirm Jevtic et al's[14] theoretical conclusions that for thermalisation with a bosonic bath, optimal temperature discrimination occurs at early times rather than in the asymptotic limit, and that for most interaction times, a thermometer initialized in a coherent superposition state outperforms one prepared in the ground state. In our case, this advantage translates to a maximum reduction of the error probability, p_e, from 47.99% to 46.19% when discriminating between temperatures $5.98\hbar\omega/k_B$ and $10\hbar\omega/k_B$ after letting the qubit interact with the bath for $t = 0.23\tau_{sp}$. This is a 90% or almost two-fold increase in the improvement over a purely random guess ($p_e = 50\%$). Furthermore, after just $t = 0.12\tau_{sp}$, the error probability is reduced, relative to a fully thermalised qubit, from 49.17% to 46.12% or approximately a 3.5 fold advantage in the improvement over random guess. We discuss the origin and limitations of this behaviour, and study the extension to the case of a thermometer composed of a finite number of spins, showing the advantages of an adaptive measurement strategy. We note that there are important differences between the optimization problem for single and multiple spins, but conclude that quantum coherence retains an advantage even in the latter case. This is a new example of a quantum metrological advantage, and may prove important for making accurate measurements of thermal properties of quantum systems with limited resources or limited disturbance.

Note: During preparation of this manuscript we became aware that similar work was being pursued by Mancino et al[30].

References

1. Giovannetti, V., Lloyd, S. & Maccone, L. Quantum-enhanced measurements: beating the standard quantum limit. *Science* **306**, 1330–1336 (2004).
2. Giovannetti, V., Lloyd, S. & Maccone, L. Advances in quantum metrology. *Nature Photonics* **5**, 222–229 (2011).
3. Aasi, J. *et al.* Enhanced sensitivity of the ligo gravitational wave detector by using squeezed states of light. *Nature Photonics* **7**, 613–619 (2013).
4. Lee, H., Kok, P. & Dowling, J. A quantum rosetta stone for interferometry. *Journal of Modern Optics* **49**, 2325–2338 (2002).
5. Mitchell, M., Lundeen, J. & Steinberg, A. Super-resolving phase measurements with a multiphoton entangled state. *Nature* **429**, 161–164 (2004).
6. Walther, P. *et al.* De broglie wavelength of a non-local four-photon state. *Nature* **429**, 158–161 (2004).
7. Nagata, T., Okamoto, R., O'Brien, J., Sasaki, K. & Takeuchi, S. Beating the standard quantum limit with four-entangled photons. *Science* **316**, 726–729 (2007).
8. Tsang, M. Quantum imaging beyond the diffraction limit by optical centroid measurements. *Physical review letters* **102**, 253601 (2009).
9. Rozema, L. *et al.* Scalable spatial superresolution using entangled photons. *Physical review letters* **112**, 223602 (2014).
10. Shin, H., Chan, K., Chang, H. & Boyd, R. Quantum spatial superresolution by optical centroid measurements. *Phys. Rev. Lett.* **107**, 083603 (2011).
11. Jaklevic, R., Lambe, J., Silver, A. & Mercereau, J. Quantum interference effects in josephson tunneling. *Physical Review Letters* **12**, 159 (1964).
12. Auzinsh, M. *et al.* Can a quantum nondemolition measurement improve the sensitivity of an atomic magnetometer? *Physical review letters* **93**, 173002 (2004).
13. Jones, J. *et al.* Magnetic field sensing beyond the standard quantum limit using 10-spin noon states. *Science* **324**, 1166–1168 (2009).
14. Jevtic, S., Newman, D., Rudolph, T. & Stace, T. Single-qubit thermometry. *Physical Review A* **91**, 012331 (2015).
15. Holevo, A. Bounds for the quantity of information transmitted by a quantum communication channel. *Problemy Peredachi Informatsii* **9**, 3–11 (1973).
16. Carmichael, H. *An open systems approach to quantum optics: lectures presented at the Université Libre de Bruxelles, October 28 to November 4, 1991*, vol. 18 (Springer Science & Business Media, 2009).
17. Mandel, L. & Wolf, E. *Optical coherence and quantum optics* (Cambridge university press, 1995).
18. Fuchs, C. Distinguishability and accessible information in quantum theory. *arXiv preprint quant-ph/9601020* (1996).
19. Nielsen, M. & Chuang, I. *Quantum computation and quantum information* (Cambridge university press, 2010).

20. Wang, D., Berry, D., de Oliveira, M. & Sanders, B. Solovay-kitaev decomposition strategy for single-qubit channels. *Physical review letters* **111**, 130504 (2013).

21. Ruskai, M., Szarek, S. & Werner, E. An analysis of completely-positive trace-preserving maps on m2. *Linear Algebra and its Applications* **347**, 159–187 (2002).

22. Lu, H. *et al.* Universal digital photonic single-qubit quantum channel simulator. *arXiv preprint arXiv:1505.02879* (2015).

23. Mohseni, M., Rezakhani, A. & Lidar, D. Quantum-process tomography: Resource analysis of different strategies. *Physical Review A* **77**, 032322 (2008).

24. Kosut, R., Walmsley, I. & Rabitz, H. Optimal experiment design for quantum state and process tomography and hamiltonian parameter estimation. *arXiv preprint quant-ph/0411093* (2004).

25. Ježek, M., Fiurášek, J. & Hradil, Z. Quantum inference of states and processes. *Physical Review A* **68**, 012305 (2003).

26. Sacchi, M. Maximum-likelihood reconstruction of completely positive maps. *Physical Review A* **63**, 054104 (2001).

27. Mahler, D. *et al.* Adaptive quantum state tomography improves accuracy quadratically. *Physical review letters* **111**, 183601 (2013).

28. Huszár, F. & Houlsby, N. Adaptive bayesian quantum tomography. *Physical Review A* **85**, 052120 (2012).

29. Higgins, B., Doherty, A., Bartlett, S., Pryde, G. & Wiseman, H. Multiple-copy state discrimination: Thinking globally, acting locally. *Physical Review A* **83**, 052314 (2011).

30. Mancino, L., Sbroscia, M., Gianani, I., Roccia, E. & Barbieri, M. Quantum simulation of single-qubit thermometry using linear optics. *To be published.* .

Acknowledgements

We thank D. Mahler and B. Sanders for useful discussions. This work was funded by NSERC, CIFAR, and Northrop-Grumman Aerospace Systems *NG Next*.

Author contributions statement

A.M.S., A.V.S., H.F., and W.K.T. jointly designed the experiment. W.K.T. performed the experiment and corresponding data analysis. A.M.S., H.F., and W.K.T. jointly conducted theoretical analysis of the qubit thermometry scheme and Bayesian adaptive state discrimination. W.K.T. wrote this manuscript along with contributions from A.M.S. and H.F.

Additional information

Competing financial interests: The authors declare no competing financial interests.

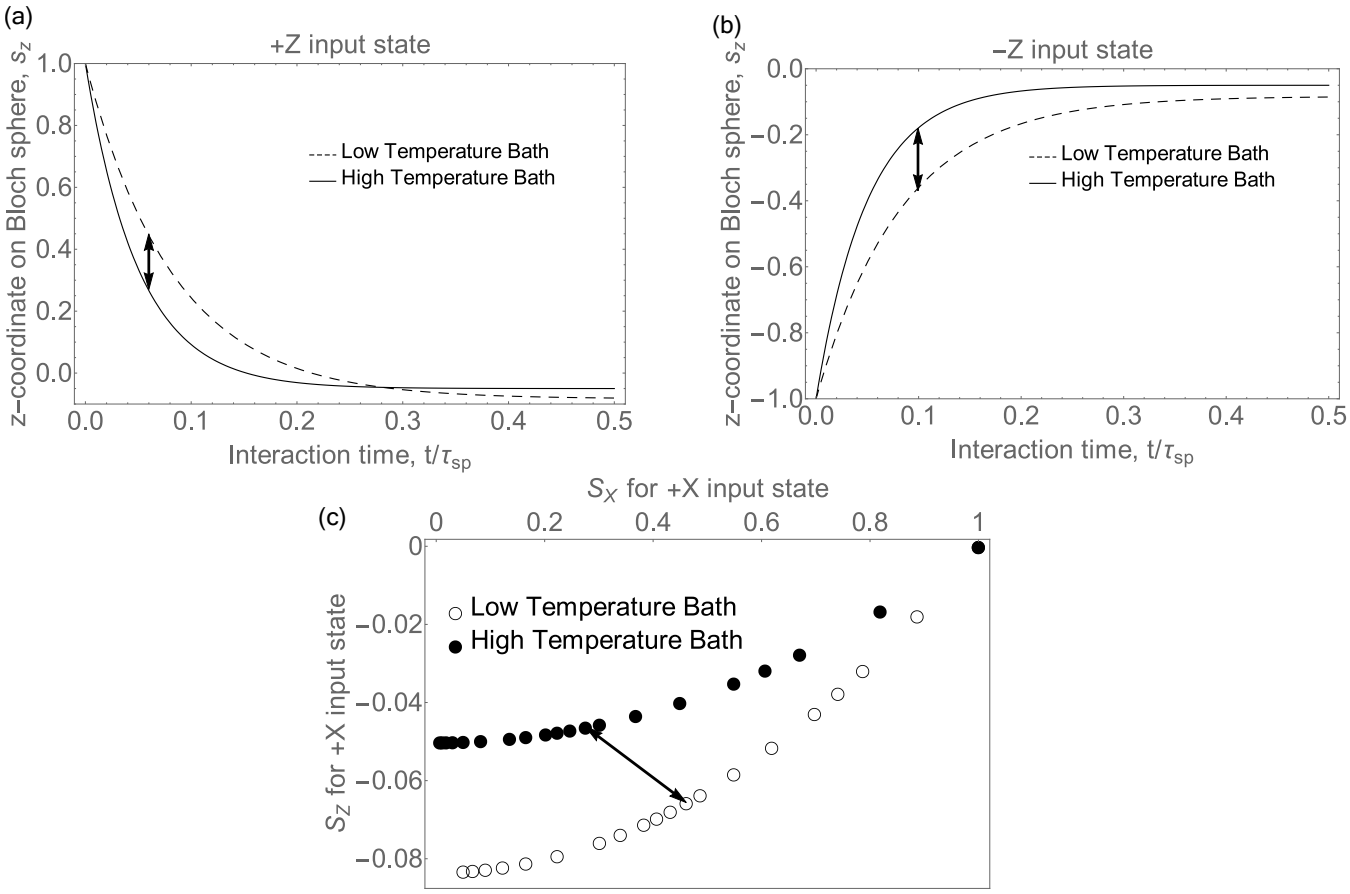

Figure 1. Bloch vector components vs interaction time. Theoretically computed components of the Bloch vector after thermalizing for t seconds, **(a)** given a $+Z$ input state, **(b)** $-Z$ input state, and **(c)** $+X$ input state. For this latter case, both s_z and s_x are shown. At $t = 0$ the state begins at the rightmost point of the bottom plot. Each subsequent timestep is shown as a pair of points, one each for high and low temperatures respectively. Arrows indicate where the greatest separation occurs.

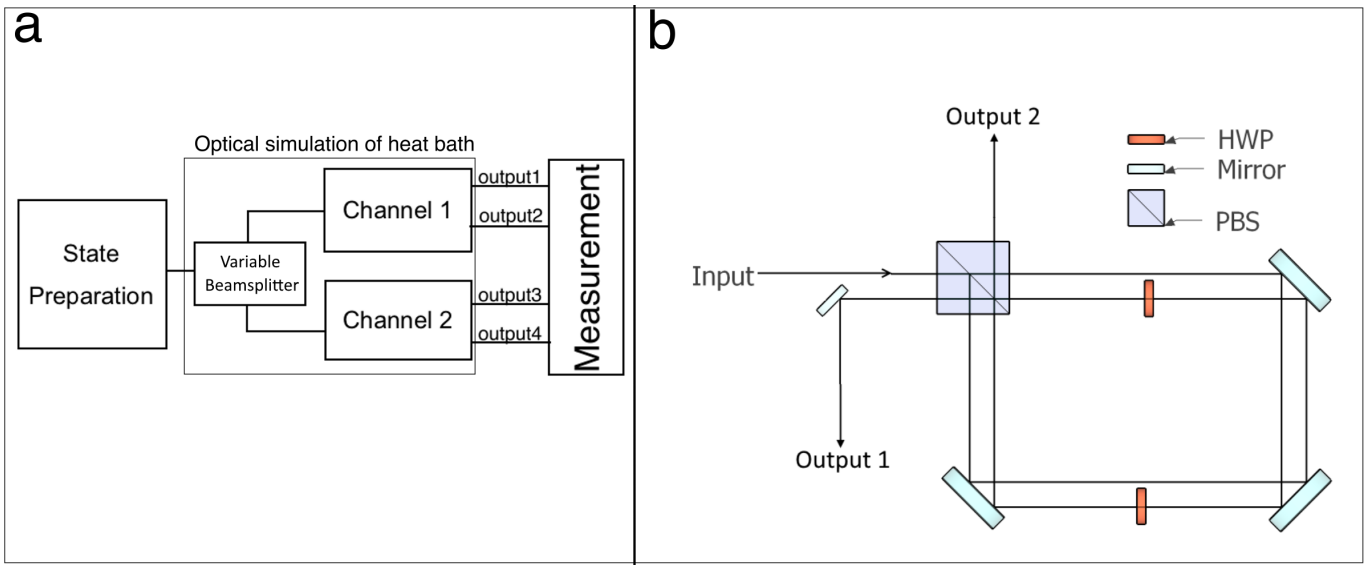

Figure 2. Experimental Scheme. a) Block schematic of the experimental setup for channel emulation. **b)** Drawing of an optical implementation of the "channels" block. Two similarly constructed copies of the channel driven by a splitter that switches between them implements a simulation of a thermal bath. The "switching" is done via a variable beamsplitter (VBS), whose coefficient of reflection/transmission can be modulated.

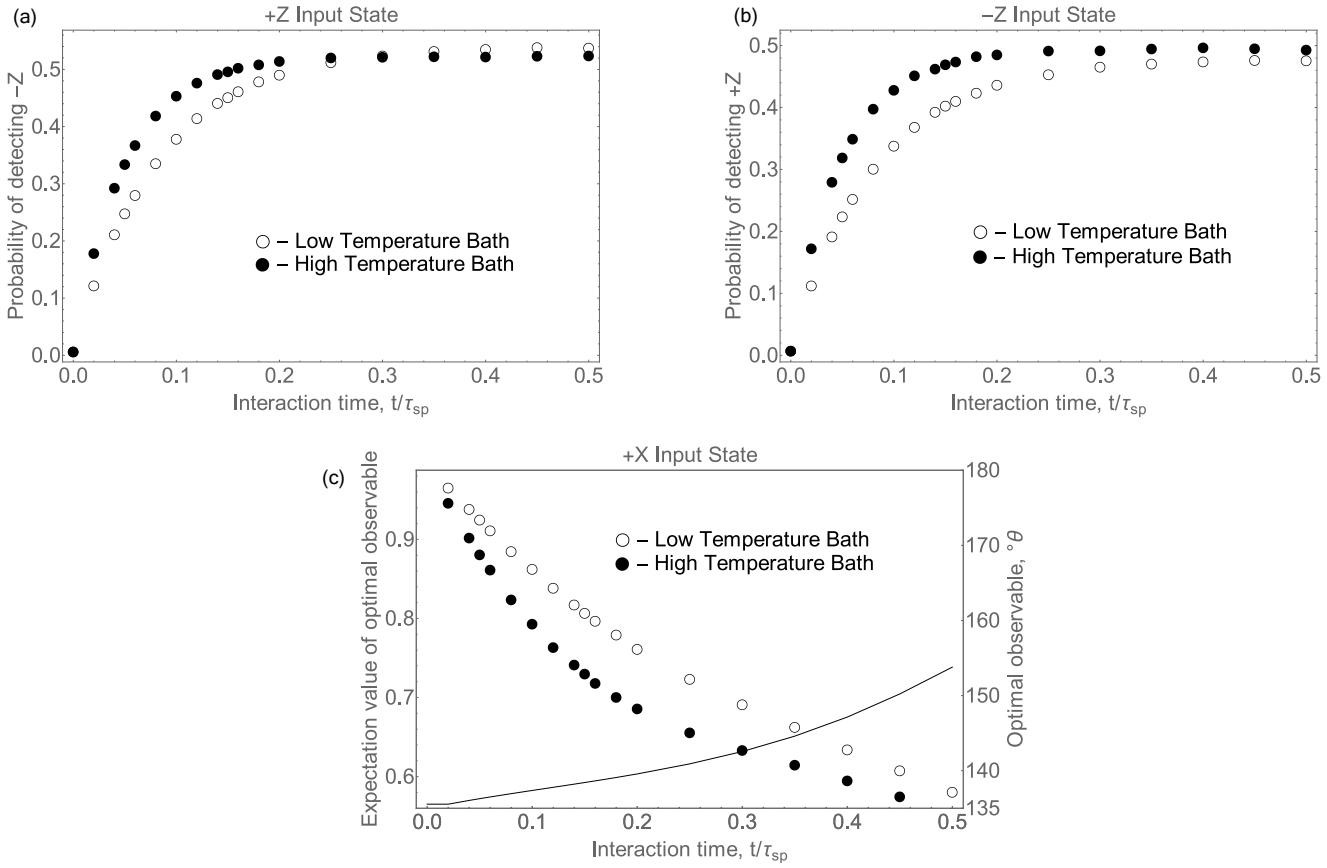

Figure 3. Plots of inferred detection probabilities. For each input state and channel setting, we computed and performed the optimal measurement for discrimination between the final states from each heat bath. **(a)** Detection probability of the $-Z$ observable (i.e. $|H\rangle\langle H|$) when input state is $+Z$ (or $|V\rangle$). **(b)** Detection probability of $+Z$ observable (i.e. $|V\rangle\langle V|$) when input state is $-Z$ (or $|H\rangle$). **(c)** Detection probability for some optimal measurement given input $+X$ (or $|D\rangle$). Although for $\pm Z$ input states the optimal observable is fixed, in the $+X$ case it varies with interaction time. In this plot, the optimal observable is parametrised as: $|\theta\rangle\langle\theta|$ where $|\theta\rangle = \cos\theta\,|H\rangle + \sin\theta\,|V\rangle$. For reference, the relaxation times for the hot bath are $\tau_2 = 2\tau_1 = 1/20 = 0.05\tau_{sp}$. For the cold bath, they are $\tau_2 = 2\tau_1 = 1/12 \approx 0.083\tau_{sp}$.

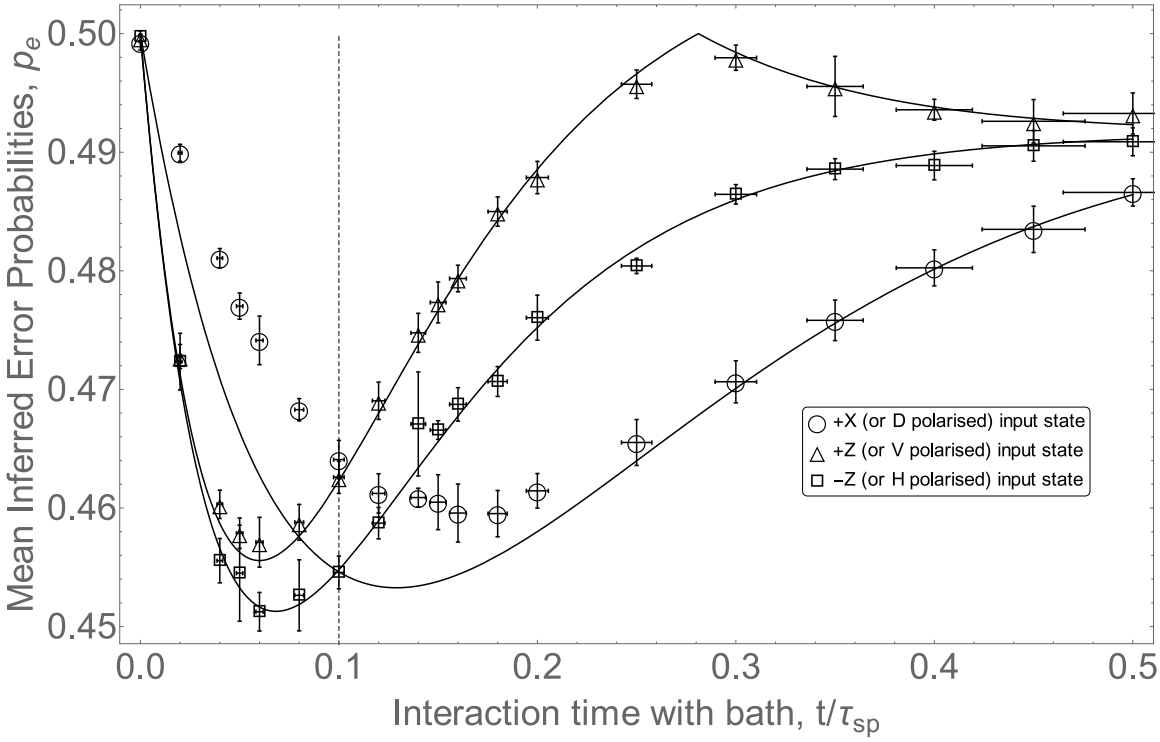

Figure 4. Inferred Error Probabilities in discriminating between final states from two heat baths. Solid lines are theory curves. Points are experimentally inferred error probabilities. Vertical error bars indicate standard deviation or spread range over multiple repetitions for a given channel setting and input state. Horizontal error bars indicate uncertainty in channel setting due to finite precision of motorised rotational mounts. Relaxation times for the hot bath are $\tau_2 = 2\tau_1 = 1/20 = 0.05\tau_{sp}$. For the cold bath, they are $\tau_2 = 2\tau_1 = 1/12 \approx 0.083\tau_{sp}$.

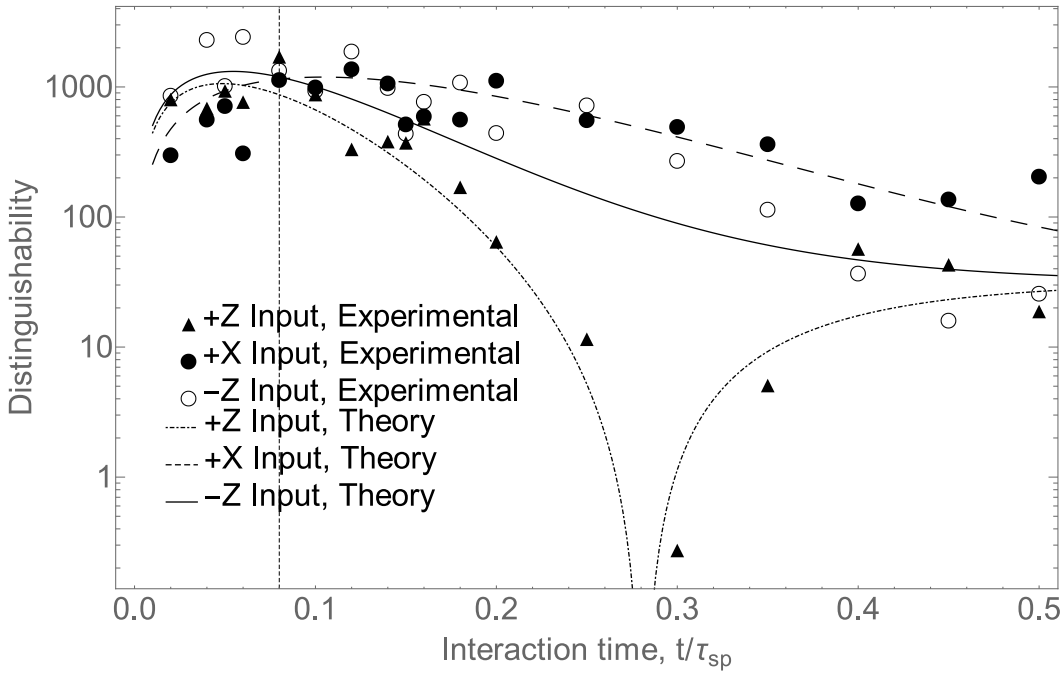

Figure 5. Plot of distinguishability in the outcome of our optimal observable (see figure 3). Higher points indicate better distinguishability. Means and variances were computed over multiple sets of comparable experimental data. Average photon number in all measurements was $\approx 40,000$.

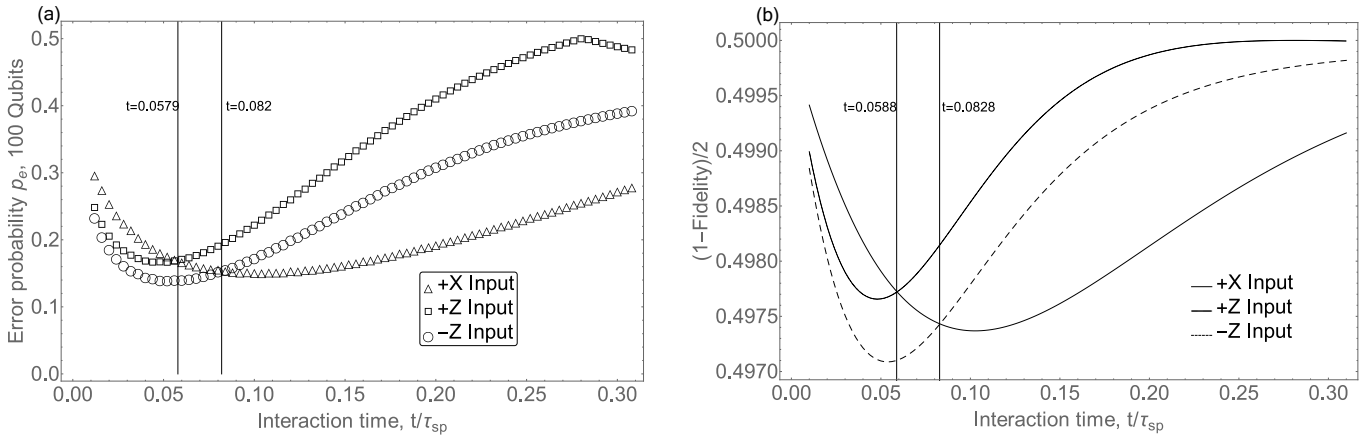

Figure 6. Performance measures for a multi-qubit thermometer. (a) Numerically computed error probabilities for 100 qubits. **(b)** The fidelity between states being discriminated (identical states have fidelity 1 whereas orthogonal states have fidelity 0).

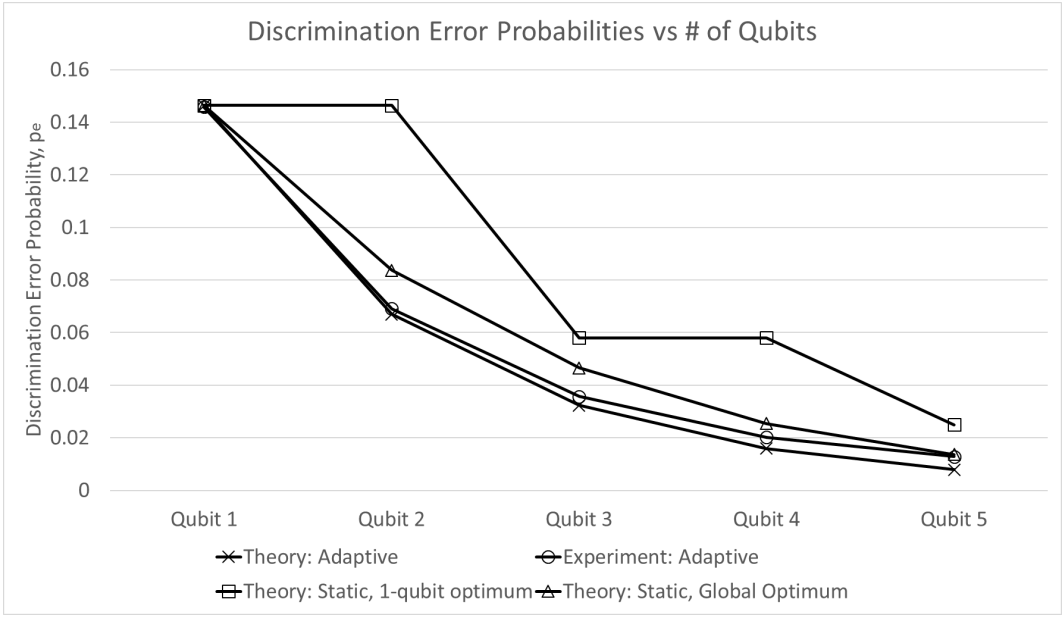

Figure 7. Plot of error probability in discriminating $|H\rangle$ from $|D\rangle$, by number of qubits for various strategies.